

Constant-degree graph expansions that preserve treewidth

Igor Markov and Yaoyun Shi

Department of Electrical Engineering and Computer Science

The University of Michigan

Ann Arbor, MI 48109-2121, USA

E-mail: {imarkov,shiyy}@eecs.umich.edu

Abstract

Many hard algorithmic problems dealing with graphs, circuits, formulas and constraints admit polynomial-time upper bounds if the underlying graph has small treewidth. The same problems often encourage reducing the maximal degree of vertices to simplify theoretical arguments or address practical concerns. Such degree reduction can be performed through a sequence of splittings of vertices, resulting in an *expansion* of the original graph. We observe that the treewidth of a graph may increase dramatically if the splittings are not performed carefully. In this context we address the following natural question: is it possible to reduce the maximum degree to a constant without substantially increasing the treewidth?

Our work answers the above question affirmatively. We prove that any simple undirected graph $G = (V, E)$ admits an expansion $G' = (V', E')$ with the maximum degree ≤ 3 and $\text{treewidth}(G') \leq \text{treewidth}(G) + 1$. Furthermore, such an expansion will have no more than $2|E| + |V|$ vertices and $3|E|$ edges; it can be computed efficiently from a tree-decomposition of G . We also construct a family of examples for which the increase by 1 in treewidth cannot be avoided.

Keywords: Treewidth, graph expansion, constant degree, ternary graph, algorithms

1 Introduction

Given a graph G , its *treewidth* $w(G)$ is a combinatorial parameter that measures to what extent a graph differs from a tree. It is defined in terms of *tree decompositions*, which are tree-like drawings of G satisfying certain constraints. The width of a specific tree decomposition represents the amount of clustering required to subsume cycles so as to make the graph look like a tree, and $w(G)$ is defined as the smallest width over all possible tree decompositions of G . Formal definitions are given in Section 2.

Since its introduction by several authors independently in the 1980's, the notion of treewidth has found numerous applications in algorithm design. Many hard combinatorial problems, such as Independent Set, Vertex Cover, SAT and #SAT, admit algorithms whose running time is $\text{poly}(n) \exp(w(G))$, where n is the input size, and G is the underlying graph structure. Thus, when $w(G) = O(\log n)$, such algorithms run in polynomial time. For example, given a CIRCUIT-SAT instance size n and a width- w tree decomposition of the circuit graph G , the *bucket elimination* algorithm [8, 6] can be used to compute the number of satisfying assignments in time $n^{O(1)} \exp(w)$. Computing the optimal tree decomposition is NP-hard [1]. But fortunately the well known algorithm by Robertson and Seymour [11] computes a tree decomposition of width $O(w(G))$ in time $|G|^{O(1)} \exp(w(G))$. Making use of this algorithm, the bucket elimination algorithm for SAT achieves the same complexity. In practice, reasonably good tree decompositions can be found by replacing the Robertson-Seymour algorithm with heuristics, which removes a runtime bottleneck and allows the entire algorithm to run fast on inputs of small treewidth, e.g., in the case of probabilistic inference [8]. The recent survey by Bodlaender [2] outlines a number of other examples.

In addition to treewidth, other graph parameters have also been heavily used in applications to estimate and moderate complexity. A particular parameter focal to this work is the maximum degree $\Delta(G)$ of a graph. It is often desirable to reduce the maximum degree of an input graph through a sequence of vertex splittings — a high-degree vertex w is replaced by an edge connecting two new vertices u and v , and each neighbor of w is assigned to be a neighbor of either u or v . Thus the original graph $G = (V, E)$ is replaced by a graph $G' = (V', E')$ that is called an *expansion* of G . The expansion process can be iterated, after which each vertex $v \in G$ is replaced by a tree T_v in G' , and each original edge $uv \in E$ corresponds to an edge in E' that connects a pair of leaves in T_u and T_v .

Degree reduction arises in several unrelated algorithmic contexts, motivated by conceptual simplification, such as the reduction of SAT to 3-SAT, or application-specific concerns. For example, VLSI circuits use gates with limited fan-in to facilitate placement and routing in dense two-dimensional silicon wafers. Large AND, OR and XOR gates frequently occur in high-level descriptions of digital logic but are routinely broken down into trees of gates with bounded fan-in. Fan-out optimization is performed because fan-outs with high electrical capacitance lead to high circuit delay. They are split into trees using buffer (repeater) gates with small fan-out. To this end, research from Intel [13] points out that the number of buffers in VLSI circuits is increasing with every technology generation and may exceed 50% of all gates in several years. To consider fan-in and fan-out optimization of VLSI circuits in our graph-based framework, we represent each gate by two vertices linked by an edge — one vertex connects all fan-ins, and the other connects all fan-outs. Performing degree minimization thus takes care of both cases.

A simple procedure (the *symmetric expansion*) constructs for a graph G an expansion G' with $\Delta(G') \leq 3$, but smaller $\Delta(G')$ cannot be guaranteed in general. This procedure replaces each vertex $u \in V$ with a path graph containing one vertex u_v for each vertex v adjacent to u in G , and replaces each $uv \in E$ by the edge connecting u_v to u_v . However, not all expansions are equally favorable. In VLSI circuit optimization, the choices of tree do not affect the circuit's functionality, but they may increase the treewidth which will complicate placement and routing.

Treewidth also features prominently in computational logic and algorithms for constraint-satisfaction

problems. In particular, *directional resolution* introduced by Davis and Putnam in 1960 for solving CNF-SAT works particularly well on instances with low width [4]. More general constraint satisfaction problems with limited treewidth admit polynomial-time algorithms for finding and counting solutions [3, 5], and the same applies to the evaluation of Bayesian networks [6]. Therefore, preserving treewidth is crucial when transforming constraint systems, e.g., when converting SAT to 3-SAT. This particular problem is the focus of [14], which proposes a specific transformation to generate 3-SAT instances at most seven times larger, whose treewidth is increased by at most one.

The considerations above lead to a natural question : is it possible to reduce the maximum degree of G through vertex splitting without substantially increasing the treewidth of G ? While treewidth cannot decrease during expansion, a simple example in Figure 1 illustrates that treewidth may dramatically increase after a symmetric expansion.

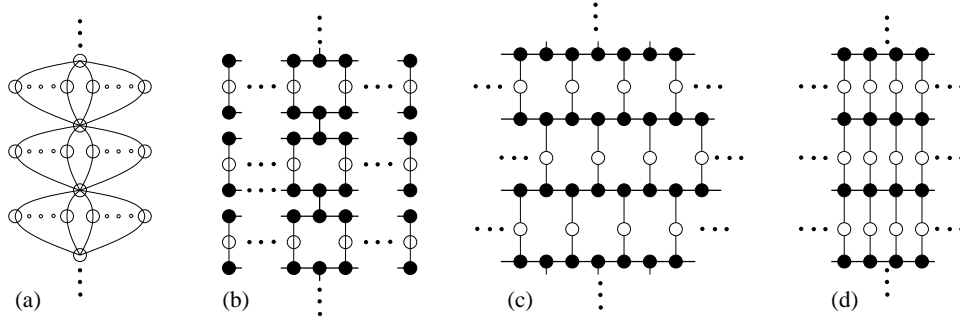


Figure 1: The graph in (a) has treewidth 2, and has a ternary expansion (b) of the same treewidth. A symmetric expansion, shown in (c), however, contains a grid minor (d). Thus (c) has treewidth at least 4 and, if scaled, can reach an arbitrary treewidth.

Let us call a graph G *ternary* if $\Delta(G) \leq 3$. Note that the maximum degree of a ternary graph cannot be further reduced by expansion. Our main result gives an affirmative answer to the above question: we prove that any graph $G = (V, E)$ admits a ternary expansion with treewidth $\leq w(G) + 1$. We give a polynomial-time algorithm to compute such an expansion from an optimal tree decomposition of G . In general, given a tree decomposition of width w , the output of the algorithm is a ternary expansion with width $\leq w + 1$. Thus, combined with the Robertson-Seymour algorithm, our algorithm outputs a ternary expansion G' of G with width $O(w(G))$ in time $|G|^{O(1)} \exp(O(w(G)))$. Finally, we construct a family of graphs G_n such that any ternary expansion G'_n of G_n must have $w(G') = w(G) + 1$. Thus our algorithm achieves the generally possible minimum treewidth. Its additional applications include a forthcoming work on quantum circuits where it helps to establish an efficient classical algorithm for simulating a broad class of quantum computations [9].

The remaining part of the paper is structured as follows. We review the notions of tree decomposition and treewidth in Section 2, then prove our main result in Section 3. Section 4 shows that our result cannot be improved, and final remarks are given in Section 5.

2 Definitions

Let $G = (V, E)$ be an undirected simple graph. For a vertex $v \in V$, we denote the set of its adjacent vertices (neighbors) by $N(v)$. Further $d(v) = |N(v)|$, and $\Delta(G) = \max_{v \in V} d(v)$.

Let G be a graph. Following definitions in [10], a *tree decomposition* of G is a tree \mathcal{T} , together with a function that maps each tree vertex w to a subset $B_w \subseteq V(G)$. These subsets B_w are called *bags* and can be used as vertex labels. In addition, the following conditions must hold.

- (T1) $\bigcup_{v \in V(\mathcal{T})} B_v = V(G)$, i.e., each vertex must appear in some bag (and may appear in multiple bags).
- (T2) $\forall \{u, v\} \in E(G), \exists w \in V(\mathcal{T}), \{u, v\} \subseteq B_w$, i.e., for each edge, there must be a bag containing both of its end vertices.
- (T3) $\forall u \in V(G)$, the set of vertices $w \in V(\mathcal{T})$ with $u \in B_w$ form a connected subtree T_u , i.e., all bags containing a given vertex must be connected in \mathcal{T} .

The *width* of a tree decomposition \mathcal{T} , denoted by $w(\mathcal{T})$, is defined by $\max_{w \in V(\mathcal{T})} |B_w| - 1$. For graph G , its treewidth $w(G)$ is the minimum width of tree decompositions of G . While NP-hard to compute in general, $w(G)$ is known for common classes of graphs [7] — a non-empty tree has treewidth 1, the $n \times n$ grid has treewidth n , and a parallel serial graph has treewidth ≤ 2 . Figure 3 shows an example of a graph of treewidth 3 and its tree decomposition of the same width.

A key motivation for the study of treewidth is the study of graph minors. Let $G = (V, E)$ be a graph. The *contraction* of an edge $uv \in E$ is the following operation on G : remove u and v (and all incident edges), and connect all neighbors of u and v to a new vertex w . A graph G' is a minor of G if G' can be obtained from a sequence of edge contractions on a subgraph of G . In this case, a tree decomposition for G also induces a tree decomposition for G' of equal or smaller width. Usually $w(G') < w(G)$, in particular, contracting all edges of G will reduce G to an empty graph, which has treewidth 0.

The process of *splitting* studied in our work can be viewed as inverse to contraction.

Definition 2.1. Let $G = (V, E)$ be a graph. The *splitting* of $v \in V$ with the support $S \subseteq N(v)$ is the following transformation of G : introduce a new vertex v' , connect v' to v , for any $s \in S$, disconnect it from v and connect it to v' . A graph G' is called an *expansion* of G if there exists a sequence of splittings that transform G to G' , and is said to be *irreducible* if no degree-2 vertex is created in the splittings. If $\Delta(G') \leq 3$, we call G' a *ternary expansion* of G .

It follows immediately from the definition of splitting a vertex v that contracting the edge vv' results in the original graph. Thus if G' is an expansion of G , then G is a minor of G' , but not *vice versa*, in general. It also follows from the definition that G' is an expansion of G if and only if G can be obtained from G' by contracting edges in a set of vertex-disjoint tree subgraphs without creating any parallel edges. Furthermore, an expansion G' of G is irreducible if and only if none of the vertices involved in the contraction has degree 2. We note that the size of any irreducible expansion must be linear in the size of the original graph. Denote by $|V|_0$ the number of degree-0 vertices in a graph $G = (V, E)$.

Proposition 2.2. Any irreducible expansion $G' = (V', E')$ of $G = (V, E)$ must have $|V'| \leq 2|E| + |V|_0$ and $|E'| \leq 3|E|$.

Proof. Let $v \in V$ be a vertex split in creating G' . Then $d(v) \geq 4$, for otherwise a degree-2 vertex would be created, contradicting to the assumption that G' is irreducible. Denote by T'_v the tree subgraph of G' whose internal vertices contract to v . Then the number of leaves of T'_v is precisely $d(v)$. Since T'_v does not have a degree-2 vertex, it follows from a simple induction that T'_v has $\leq d(v) - 2$ internal vertices and $\leq d(v) - 1$ internal edges. Therefore,

$$|V'| \leq \sum_{v \in V: d(v) \geq 4} (d(v) - 2) + \sum_{v \in V: d(v) \leq 3} 1 \leq \sum_{v \in V} d(v) + |V|_0 \leq 2|E| + |V|_0,$$

and

$$|E'| \leq |E| + \sum_{v \in V: d(v) \geq 4} (d(v) - 1) \leq 3|E|.$$

□

In our construction of an expansion, we may introduce degree-2 vertices for the convenience of bounding the treewidth. Such vertices can be removed easily at the end to obtain an irreducible expansion.

3 Main result and its proof

Theorem 3.1. *There is a polynomial-time algorithm that, given a graph $G = (V, E)$ and its tree decomposition of width w , computes a ternary expansion $G' = (V', E')$ with $w(G') \leq w + 1$. In particular, G admits a ternary expansion whose treewidth is no more than $w(G) + 1$.*

The construction of G' takes several stages: first we construct graph G_1 with a tree decomposition \mathcal{T}_1 such that the subgraph of G_1 induced by each bag in \mathcal{T}_1 has maximum degree ≤ 2 . In the second stage, each vertex v is split many times to replicate the structure T_v , the tree formed by those bags containing v in \mathcal{T}_1 . Two vertex trees T_u and T_v for $uv \in E(G_1)$ are then connected through a pair of vertices corresponding to the same bag that contains u and v . In the last stage, each vertex is split many times to reduce the degree within its vertex tree. We combined the last two stages in our following description of the construction.

Proof of Theorem 3.1. If $w(G) \leq 1$, then G is a forest. Repeatedly splitting a vertex with degree ≥ 4 with 2 supporting vertices will result in an expansion $G' = (V', E')$ with $\Delta(G') \leq 3$, $|V'| \leq |V| + |E| = 2|V| - 1$, and $w(G') \leq 1$. Thus the Theorem holds. We now consider $w(G) \geq 2$.

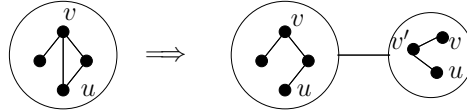


Figure 2: Inside a bag, vertex v of degree 3 is split with the support of a neighbor u in the same bag. A new neighboring bag is created, containing v , u , and the new vertex v' from the splitting.

Stage 1: Reducing the maximum degree inside a bag. We sequentially scan the bags of \mathcal{T} and apply the following operations. If a bag B contains a vertex v with $d(v) \geq 3$, let $u \in B$ be a neighbor of v , split v with the support $\{u\}$. Denote the new vertex by v' . (This is equivalent to placing v' at the edge uv .) Create a new bag B' containing $\{u, v, v'\}$ and attach it to B . This process extends the u -subtree and the v -subtree of \mathcal{T} by a leaf bag B' , and adds a trivial v' -subtree (B'), thus results in a tree decomposition. Since the new bag has 3 vertices, and $w(\mathcal{T}) \geq 2$, the width of the new tree decomposition remains the same. Figure 2 illustrates the process of adding one bag.

Denote by $G_1 = (V_1, E_1)$ the resulting graph, and by $\mathcal{T}_1 = (\mathcal{V}_1, \mathcal{E}_1)$ the resulting tree decomposition. Let $k \leq |E|$ be the number of splittings. We have

(3.a) The maximum degree of an induced subgraph of G_1 by vertices in a bag of \mathcal{T}_1 is at most 2.

(3.b) $w(\mathcal{T}_1) = w(\mathcal{T})$.

Stage 2: Completing the construction. We now construct a graph $G' = (V', E')$ from G_1 and later show that it is an expansion of G_1 , thus an expansion of G . Fix a total ordering \preceq on \mathcal{V}_1 . Define $V' \subseteq V_1 \times \mathcal{V}_1 \times \mathcal{V}_1$ as follows

$$V' = \{(v, B, B') : (v \in B \cap B') \wedge (B = B' \vee BB' \in \mathcal{E}_1)\}.$$

For each $v \in V_1$, let T_v be the v -subtree of \mathcal{T}_1 . There are three types of edges in G' .

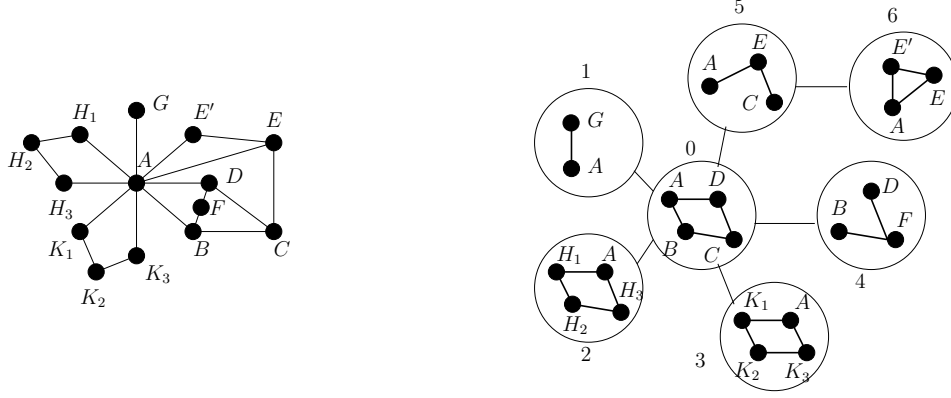


Figure 3: A graph and a tree decomposition that satisfies the condition in Stage 2 (of having ≤ 2 induced degree within each bag).

- (3.i) Let B_0 be a bag, $v \in B_0$, and $B_1 \prec B_2 \prec \dots \prec B_t$ are the neighbors of B_0 in T_v . An *ordering edge* connects a vertex (v, B_0, B_i) with (v, B_0, B_{i+1}) , $0 \leq i \leq t-1$.
- (3.ii) An *intra-tree edge* connects (v, B, B') with (v, B', B) , where $B \neq B'$, and $BB' \in E(T_v)$. In this case, $v \in B \cap B'$.
- (3.iii) An *inter-tree edge* connects (v, B, B) with (u, B, B) , for $u \neq v$, $uv \subseteq E_1$. In this case, $u, v \in B$.

The edge set E' consists of all possible ordering and intra-tree edges, but only one inter-tree edge for one edge of E_1 . Figures 3 and 5 illustrate the operations in Stage 2.

Stages 1 and Stage 2 take polynomial time. Using the following claims, we verify that G' satisfies the properties in the Theorem. \square

Claim 3.2. The graph G' is an expansion of G .

Proof. Contracting all ordering edges associated with an occurrence of v in a bag B_0 combines all (v, B_0) into a single vertex associated with this occurrence. All those vertices are connected through intra-tree edges, and this graph is precisely T_v . Since T_v and T_u are connected through one inter-tree edge if and only if $uv \in E$, contracting the intra-tree edges (after contracting the ordering edges) gives G . Since no parallel edges were created in the whole process, G' must be an expansion of G . \square

Claim 3.3. The graph G' satisfies $\Delta(G') \leq 3$.

Proof. A vertex (u, B, B) is incident to at most one ordering edge, no intra-tree edge, and at most two inter-tree edges. By Property (3.a), its degree ≤ 3 . A vertex (u, B, B') with $B \neq B'$ is incident to at most two ordering edges, at most one intra-tree edge (connecting (u, B', B)), and no inter-tree edge. Therefore its degree ≤ 3 , and $\Delta(G') \leq 3$. \square

Claim 3.4. The graph G' satisfies $w(G') \leq w + 1$.

Proof. Our construction of a tree decomposition \mathcal{T}' for G' uses the following procedure. Suppose at some stage we have constructed two bags represented as ordered sets $A = \{v_1, v_2, \dots, v_t\}$, $B = \{v'_1, v'_2, \dots, v'_t\}$. Then *bridging* the two bags involves the following steps:

(Step 1) Create bags B_1, B_2, \dots, B_t , where $\forall i, 1 \leq i \leq t$,

$$B_i = \{v_1, v_2, \dots, v_{i-1}, v_i, v'_i, v'_{i+1}, \dots, v'_t\}.$$

(Step 2) Connect A to B_1 , B_i to B_{i+1} , $1 \leq i \leq t-1$, and B_t to B .

We refer to the constructed path graph as the *bridge* bridging A and B , and each B_i as a vertex in the bridge. Note that for each i , $1 \leq i \leq t$, $|B_i| \leq t+1$, and v_i or v'_i appear together in B_i . When we apply bridging, vertices in the two bags will be ordered so that the pair v_i and v'_i take values (v, B_1, B_2) and (v, B'_1, B'_2) in V' for some common underlying $v \in V_1$.

We now describe the construction of \mathcal{T}' .

(Step 1) For each bag $B_0 \in \mathcal{V}_1$, adjacent to bags $B_1 \prec B_2 \prec \dots \prec B_t$ in \mathcal{T}_1 , create bags (B_0, B_i) , $\forall i, 0 \leq i \leq t$. For each $v \in B_0$, let $B_{i_1} \prec B_{i_2} \prec \dots \prec B_{i_\ell}$ be the neighbors of B_0 in T_v . Place (v, B_0, B_0) in (B_0, B_0) (thus an inter-tree edge connecting (v, B_0, B_0) and (u, B_0, B_0) is contained in the bag (B_0, B_0)). For all j , $1 \leq j \leq \ell$, place (v, B_0, B_{i_j}) in (B_0, B_i) , $\forall i, i_{j-1} < i \leq i_j$, with $i_0 = 0$. In particular, any $(v, B, B') \in V'$ is contained in (B, B') .

(Step 2) Bridge (B_0, B_i) with (B_0, B_{i+1}) , $0 \leq i \leq t-1$. Thus the ordering edge connecting (v, B_0, B_i) and (v, B_0, B_{i+1}) is contained in some vertex of this bridge.

(Step 3) If $BB' \in \mathcal{E}_1$, bridge (B, B') and (B', B) . Thus an intra-tree edge connecting (v, B, B') and (v, B', B) is contained in a vertex of this bridge.

Figure 6 shows part of the tree decompositions before the bridging operations for the ternary expansion in Figure 5.

We have shown above that each edge is contained in some bag. Thus we need only to show that bags containing a vertex must form a subtree. By construction, any vertex (v, B, B) is contained in a single bag (B, B) , which forms a tree. Let $(v, B_0, B) \in V'$ and $B_0 \neq B$. Suppose the neighbors of B_0 in \mathcal{T}_1 are $B_1 \prec B_2 \prec \dots \prec B_t$, and those containing v are $B_{i_1} \prec B_{i_2} \prec \dots \prec B_{i_\ell}$. Also, $B = B_{i_j}$, for some j , $1 \leq j \leq \ell$, and let $j_0 = i_{j-1} + 1$ (with $i_0 = 0$). Then (v, B_0, B) is placed in (B_0, B_{j_0}) , (B_0, B_{j_0+1}) , \dots , (B_0, B_{i_j}) , as well as all the vertices bridging them together. In addition, (v, B_0, B) also appears in a connected subgraph of the bridge bridging (B_0, B) and (B, B_0) . Thus the bags containing (v, B_0, B) form a path. This completes the proof that \mathcal{T}' is a tree decomposition for G' .

Recall that (B, B) and B are of the same size, while for $i \neq 0$ $|(B_0, B_i)| \leq |B|$. Furthermore, the size of a vertex in a bridge is at most 1+ the size of the end bag. Therefore, we conclude that $w(\mathcal{T}') \leq w(\mathcal{T}_1) + 1 = w(\mathcal{T}) + 1 = w + 1$. \square

4 The treewidth bound in Theorem 3.1 is sharp

We now demonstrate a family of graphs of growing treewidth, for which any ternary expansion increases treewidth by one.

Let K_n be the complete graph with n vertices (n -clique), and $K_{n,n}$ be the complete bipartite graph whose two sets of vertices are denoted by $\{v_i : i \in \mathbb{Z}_n\}$ and $\{v'_i : i \in \mathbb{Z}_n\}$. The graph $\bar{K}_{n,n}$ is obtained from $K_{n,n}$ by deleting the edges $v_i v'_i$, $i \in \mathbb{Z}_n$. Recall that for a graph with $2n$ vertices, a *perfect matching* is a subgraph consisting of n vertex-disjoint edges. We call two edges $e_1, e_2 \in E$ *engaged* if there is an edge $uv \in E$ such that e_1 and e_2 are incident to u and v , respectively.

Definition 4.1. Let $G = (V, E)$ be a graph with $|V|$ even. A perfect matching M of G is called a *bramble matching* if any two edges in M are engaged in G .

Since contracting edges in a bramble matching of $G = (V, E)$ produces $K_{|V|/2}$, we immediately obtain

Proposition 4.2. If $G = (V, E)$ admits a bramble matching then $w(G) \geq |V|/2 - 1$.

Proposition 4.3. For $n \geq 3$, $w(\bar{K}_{n,n}) = n - 1$.

Proof. If $n \geq 3$, $i - 1 \neq i + 1$, for any $i \in \mathbb{Z}_n$. Thus the edges $v'_{i-1}v_i$ and v'_iv_{i+1} are connected by the edge $v'_{i-1}v_{i+1}$. Therefore, the edges $\{v_iv'_{i+1} : i \in \mathbb{Z}_n\}$ form a bramble matching. By Proposition 4.2, $w(\bar{K}_{n,n}) \geq n - 1$. On the other hand, the following is a tree decomposition of $\bar{K}_{n,n}$ with width $n - 1$: a center bag containing $\{v_i : i \in \mathbb{Z}_n\}$, and n bags B_k , $k \in \mathbb{Z}_n$, connected to it, where B_k includes $\{v_i : i \neq k\} \cup \{v'_k\}$. Thus $w(\bar{K}_{n,n}) = n - 1$. \square

Proposition 4.4. Let $n \geq 19$, and G_n be a ternary expansion of $\bar{K}_{n,n}$. Then $w(G_n) \geq n$.

Proof. For a vertex $v \in V(\bar{K}_{n,n})$, denote its expansion in G_n by T_v (which is a tree). Contract all vertices in T_{v_i} and $T_{v'_i}$ for $i \neq 0$. According to the Tree-partitioning Lemma 5.1 in Appendix B, there must exist an edge uv in T_{v_0} such that the number of leaves closer to u (than to v) is between $(n-1)/3$ and $2(n-1)/3$. Contract all edges other than uw . Similarly, in $T_{v'_0}$, contract all edges except an edge $u'w'$ having the same property as uw . Denote the resulting graph $G = (V, E)$. Then each of $\{u, w, u', w'\}$ is connected to $\geq (n-1)/3 \geq 6$ neighbors from the rest of V . Without loss of generality, assume that the following are edges of G :

$$\{uv'_1, wv'_2, uv'_3, wv'_4, u'v_5, w'v_6, u'v_7, w'v_8\}.$$

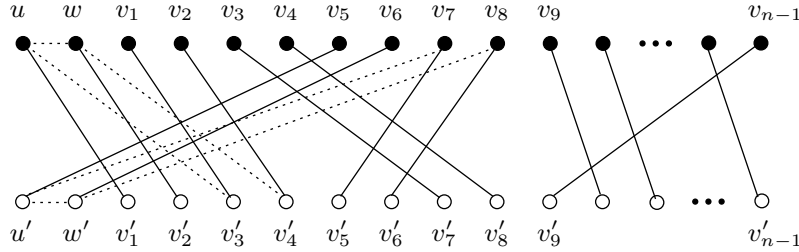


Figure 4: The solid lines form the bramble matching M , defined in Equation 1.

We construct a bramble matching M of $n + 1$ edges as shown in Figure 4: $M = M_1 \cup M'_1 \cup M_2$, where

$$\begin{aligned} M_1 &= \{uv'_1, wv'_2, u'v_5, w'v_6\}, \\ M'_1 &= \{v_1v'_3, v_2v'_4, v_3v'_7, v_4v'_8, v'_5v_7, v'_6v_8\}, \quad \text{and,} \\ M_2 &= \{v_iv'_{i+1} : 9 \leq i \leq n-2\} \cup \{v_9v_{n-1}\}. \end{aligned} \tag{1}$$

By direct inspection, M is a bramble matching. Therefore, $w(\bar{K}_{n,n}) \geq n$, by Proposition 4.2. \square

5 Discussion

The main result of this paper is good news for many application domains and generally means that sparsification of instances does not adversely affect their intrinsic complexity, if performed carefully. Theorem

3.1 shows that such an expansion does not need to increase the treewidth by more than 1. For example, extending CNF-SAT with equality clauses (which is trivially supported by most SAT solvers today), one can reduce CNF-SAT to 3,3-SAT, where each clause has up to three literals and every variable participates in at most clauses. Such a reduction could simplify the design and implementation of high-performance SAT solvers, which have become a popular topic in the last 10 years, thanks to important applications in AI, VLSI CAD, logistics, etc. In this special case, our graph-theoretical result is consistent with domain-specific work in [14], which reduces SAT to 3-SAT in a different way, but may also increase treewidth by at most one. Another application would be to optimize VLSI circuits for better layout by breaking down large AND, OR, XOR gates into trees of smaller gates and applying fan-out optimization using buffer insertion (as outlined in the Introduction).

The step-by-step description of our algorithm in the proof of Theorem 3.1 may seem daunting, but the main insight behind this algorithm is rather simple — to expand a given vertex, one must replicate the tree structure of a good tree decomposition around this vertex. We hope that the idea to reuse the local structure of tree decompositions will find other uses as well. Perhaps, the most surprising part of our work is the detailed analysis of how treewidth can change during the proposed construction — it can grow only by 1, regardless of the parameters of the input graph, and sometimes cannot be preserved by any expansion.

An interesting direction to extend our main result is to avoid the reliance of our algorithm on a tree decomposition. Perhaps, such an algorithm might also help in constructing a tree decomposition or determining the treewidth. It is also an interesting graph-theoretical problem to characterize the class of graphs that admit ternary expansions of the same treewidth. Not containing the graphs in our example as a minor appears to us a likely characterization.

References

- [1] S. Arnborg, D. G. Corneil, and A. Proskurowski. Complexity of finding embeddings in a k -tree. *SIAM J. Algebraic Discrete Methods*, 8(2):277–284, 1987.
- [2] H. L. Bodlaender. Treewidth: Characterizations, applications, and computations. Technical Report UU-CS-2006-041, Institute of Information and Computing Sciences, Utrecht University, 2006.
- [3] R. Dechter and J. Pearl, Network-based Heuristics for Constraint-Satisfaction Problems, *Artificial Intelligence*, 34(1):1–38, 1987.
- [4] R. Dechter and I. Rish, Directional Resolution: The Davis-Putnam Procedure Revisited, *Proc. Principles of Knowledge Representation and Reasoning (KR)*, 134–145, 1994.
- [5] R. Dechter, K. Kask, E. Bin, and R. Emek. Generating Random Solutions for Constraint-Satisfaction Problems *Proc. AAAI/IAAI* 15–21, 2002.
- [6] R. Dechter. Bucket elimination: a unifying framework for reasoning. *Artificial Intelligence*, 113(1-2):41–85, 1999.
- [7] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, third edition, 2005.
- [8] S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *J. Roy. Statist. Soc. Ser. B*, 50(2):157–224, 1988.

- [9] I. Markov and Y. Shi. Simulating quantum computation by contracting tensor networks. To appear in *SIAM Journal on Computing*, 2007.
- [10] N. Robertson and P. D. Seymour. Graph minors. III. Planar tree-width. *J. Combin. Theory Ser. B*, 36(1):49–64, 1984.
- [11] N. Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52(2):153–190, 1991.
- [12] R. Seidel. A new method for solving constraint satisfaction problems. In *International Joint Conference on Artificial Intelligence (IJCAI-81)*, 338–342, 1981.
- [13] P. Saxena, N. Menezes, P. Cocchini and D.A. Kirkpatrick. Repeater scaling and its impact on CAD. *IEEE Trans. on Computer-Aided Design of Integrated Circuits*, 23(4):451–463, 2004.
- [14] Y. Zabiyaka and A. Darwiche, On tractability and hypertree width, *Technical Report D-150*, UCLA Computer Science, 2007.

Appendix A: Tree-partitioning Lemma

Lemma 5.1. *Consider a tree T with $k \geq 3$ leaves in which no vertex has degree greater than three. While every tree-edge separates the tree into two connected components, for at least one edge each component contains between $k/3$ and $2k/3 - 1$ leaves, inclusively.*

Proof. We regard T as a tree rooted at a non-leaf vertex r . Note that each vertex, except r , has no more than two children. For a vertex v of T , define the size of v , $s(v)$, to be the number of leaves of the subtree rooted at v . Let v be the child of r with the largest size among its siblings. Then $s(v) \geq k/3$. If $s(v) < 2k/3$, then rv satisfies the requirement. Otherwise, traverse down the tree following the edge that connects the child with the larger size. Suppose uv is the edge so that v is the first vertex encountered having $< 2k/3$ size. Since $s(u) \geq 2k/3$, and $s(v) \geq s(u)/2 \geq k/3$, uv satisfies the requirement. \square

Appendix B: Figures

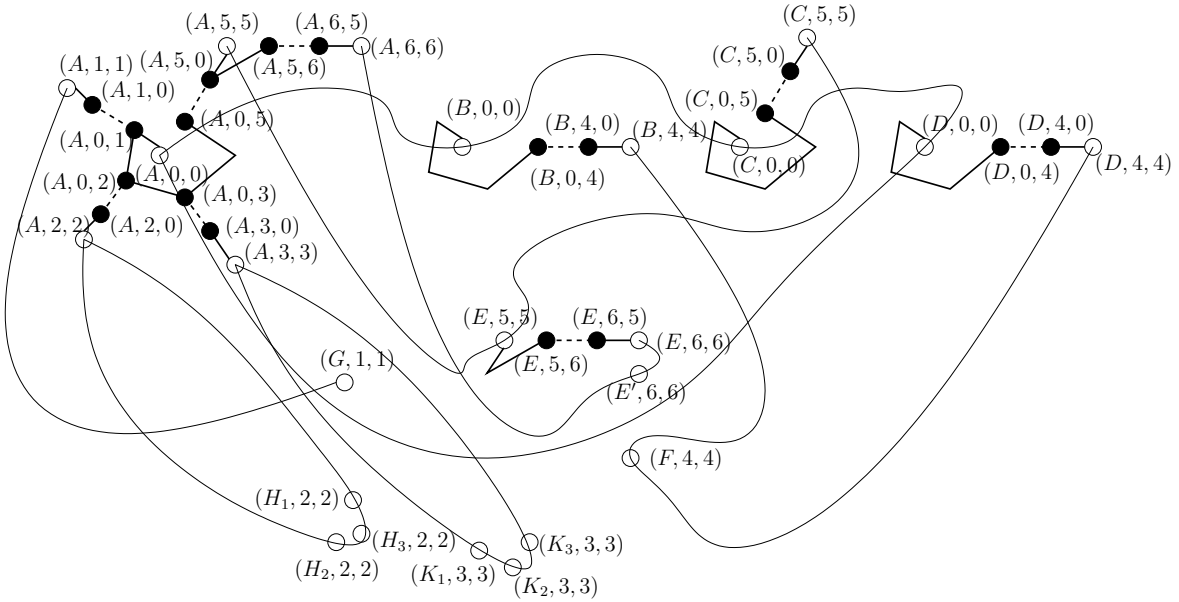


Figure 5: The ternary expansion of the graph in Figure 3 after the steps in Stage 2. The straight solid lines are *ordering* edges, the dashed lines are *intra-tree* edges, and the curves are *inter-tree* edges. The expansions of vertex A , B , C , and D are drawn in separate positions at the upper part, and all the others are at the lower part.

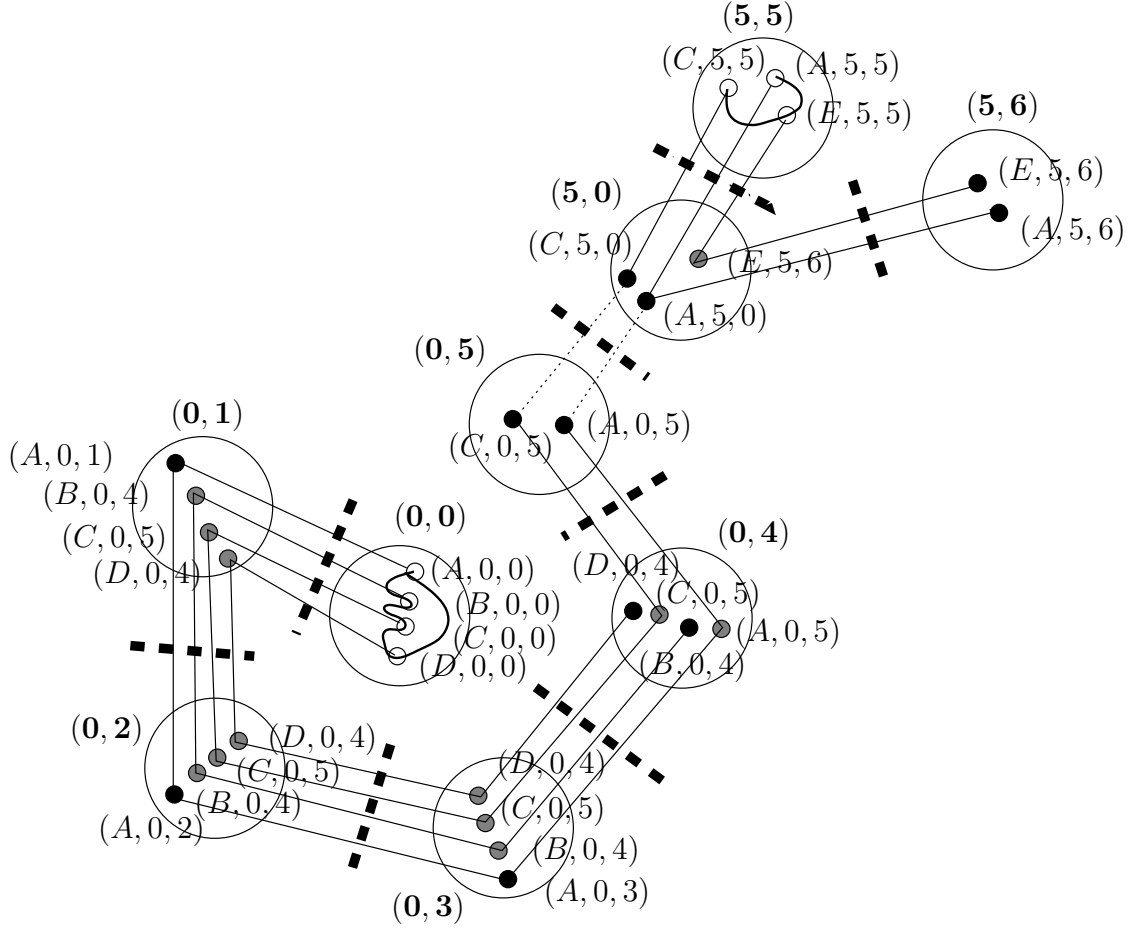


Figure 6: Part of the final tree decomposition corresponding to the original Bags 0 and 5 *before* the bridging operations, which will create a sequence of new bag between bags separated by dash bars in the Figure. The shaded vertices are the re-occurrences of a non-shaded (either black or white) vertices. The edges between non-shaded vertices will be contained in some are contained in some bag. The bridging operations will increase the treewidth by 1.